



# Control Plane

**The Internet is not deterministic; it is a best effort system with neither guarantees nor perfection.**

**Anyone who says they can deliver either is selling something ... or getting ready to fail**

Colonel Tim Gibson, Ph.D.  
DARPA Advanced Technology Office  
tgibson@darpa.mil  
703.526.4764



# A New Network Paradigm



- Change how we use the network—not how it works

**IP Networks are:**

- Best Effort
- Not deterministic

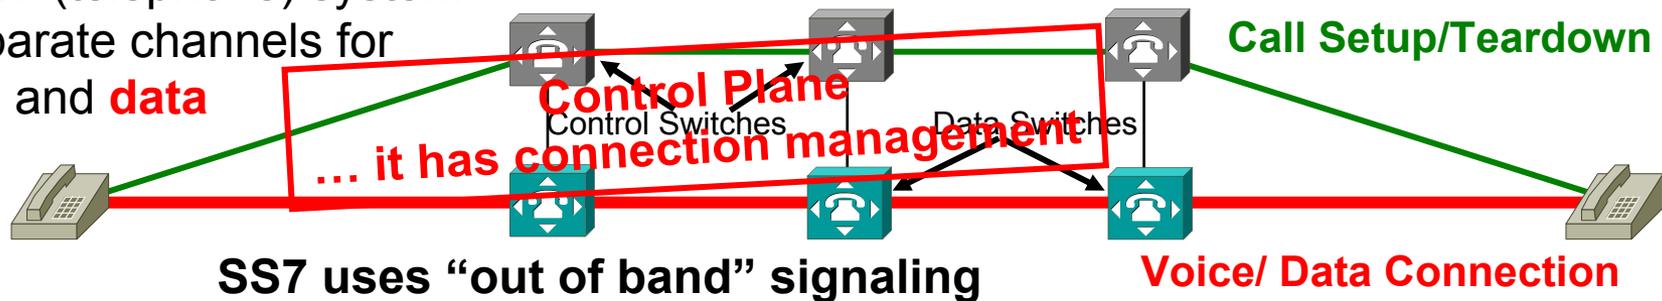
- Stop trying to impose determinism and guarantees

**Don't confuse quality of service with traffic engineering**

- Make the best effort 'the best it can possibly be'

**Routing paths, path selection, transmission characteristics**

The SS7 (telephone) system has separate channels for **control** and **data**



Network Control Plane: A mechanism connection management devices use to control and access network components and services

Data and routing information is **separate**

Advantages:

- Control isolated from users
- Guaranteed quality upon connection
- Attribution of end points

Disadvantages:

- Expensive infrastructure
- Inflexible architecture
- Difficult to add nodes (telephone number!)
- Limited services

TCP/UDP/ICMP in IP packets  
Combines data and routing information



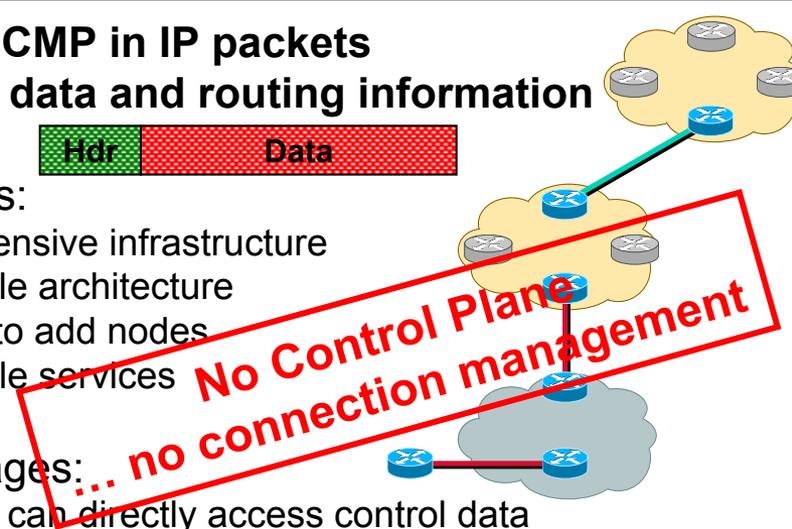
Advantages:

- Inexpensive infrastructure
- Flexible architecture
- Easy to add nodes
- Flexible services

Disadvantages:

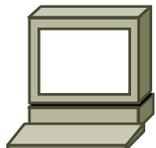
- Users can directly access control data
- No guaranteed connection quality
- No end point attribution

IP uses "in band" signaling





# What we have today *instead* of a network with a control plane



Hosts are the *de facto* connection management devices

## ICMP (Internet Control Message Protocol)



## RSVP and QoS

*The network as an oracle*

SNMP (Simple Network Management Protocol): **Not accessible** to normal hosts

Network as a 'black box'

- Transmit if they can seize the local media channel
- No knowledge of link to destination
- Focus on packet and flow control
- Provides the host no information about the network state
- Not useful to applications
- Add-ons to reserve capacity or guarantee service levels
- Not widely implemented because of infrastructure changes
- Asking network for something when ignorant of network state
- Bits go in and usually get delivered
- If not delivered, generate an error
- Hosts must infer everything about the network

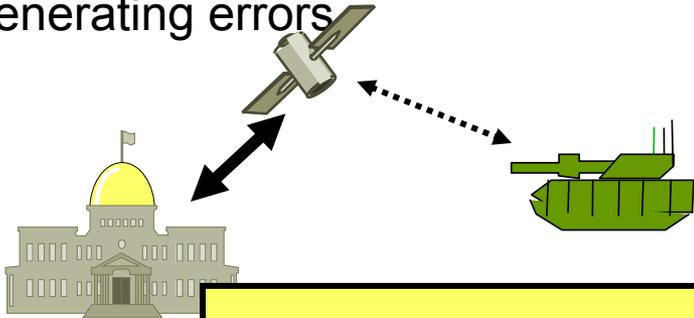
Hosts—connection managers—should **know** the network state ... not infer it!



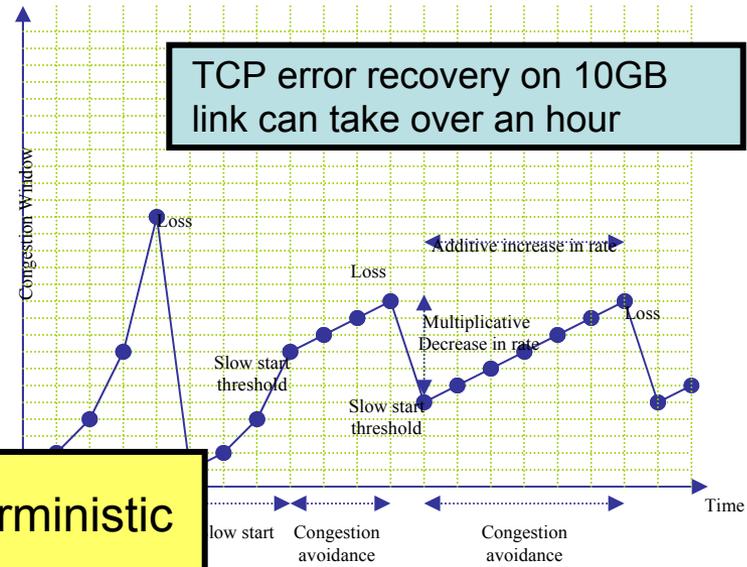
# 'Black Box' Networks: Impact on Defense Department Systems



High capacity, high quality CONUS links feed into low capacity, long latency, poor quality tactical links → flooding tactical links and generating errors

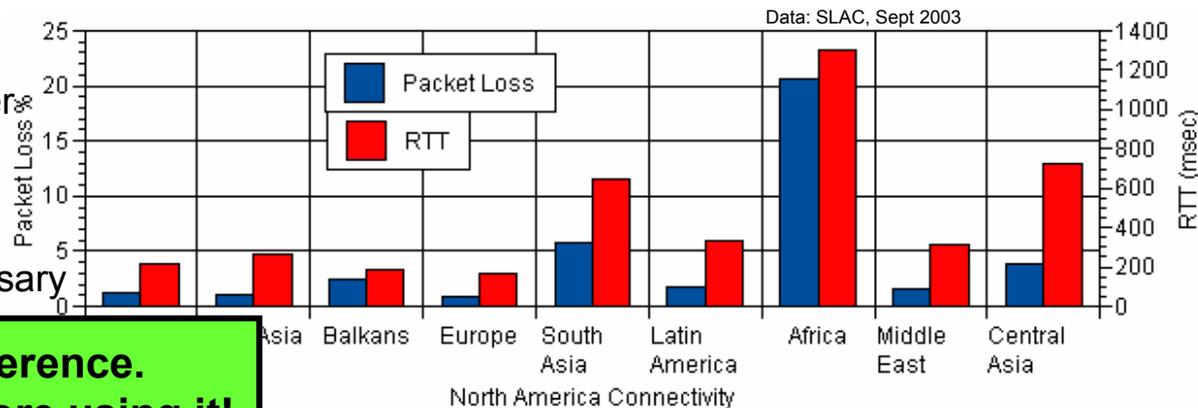


**Need to make the network more deterministic**



Results in NCW disaster:

- 40% congestion common
- Orders transmissions take longer
- Real-time displays aren't
- Common practice is to increase link-speed at 50% utilization
- Uses more satellites than necessary



**Stop relying on black box inference.  
Find out the network's state before using it!**

**Poor link quality and long latency in deployment areas**





# Program Goals



## Working system that:

- Improves path efficiency
  - Total bits transmitted to useful data received on same link
  - Network predictability (*i.e.*, more deterministic)
    - Given relevant network performance, accurately forecast near-term expected performance
- Provides choice between different connection qualities
  - Given connection diversity, choose the best quality connection
    - “Quality” is customer driven by delay, loss, jitter, throughput and packet fragmentation
- Pushes unwanted traffic off networks
  - Given unwanted internal network traffic exists, remove it
    - “Push” initiated by the user/host/connection management device



# Metric Definitions



- End-to-end Throughput
  - Measured by the amount of time it takes to transfer a given data set from one user to another
    - “Pulse” Test – 10-30 sites (users) transferring 1-10 GB simultaneously
      - Variable connection quality
    - Connectivity Test – 2-5 sites transfer up to 10 GB of various sizes
      - Multiple connections paths, each with varying quality
- Improvement Costs
  - The amount of additional funds the new system will cost.
    - Given the initial cost of a Planet-Lab installation, installing any improvements will ideally cost no more than an additional 10%
    - Scales to entire Defense Department network



# Go / No Go Metrics



	Today	Milestone 1	Milestone 2	Milestone 3
End-to-end throughput	--	3x	6x	10x
Improvement Costs	--	+30%	+20%	+10%

Sanity Test: Convene board of experts (DARPA and others) to review test results and expected infrastructure improvement costs at the end of Milestones One and Two



# Performance Testing



## Planet Lab

- 150+ Institutions World-Wide
- WAN within the Internet
- Tunneled connections
- Fully instrumented
- Allows WAN testing and measurement
- Require performers to join



- Saves expense of establishing infrastructure (\$2M+)
- Provides experimentation platform
- Tests on real networks

### Phase 1

- Planet-Lab performance tests over several weeks

### Phase 2

- Planet-Lab performance tests for three weeks
- Limited NIPRNET testing with a Service or Combatant Command

### Phase 3

- Focused test during a major exercise
  - Army, Navy, Combatant Command
- Test CONUS to 'foxhole' links





# Performance Testing

## Milestone One



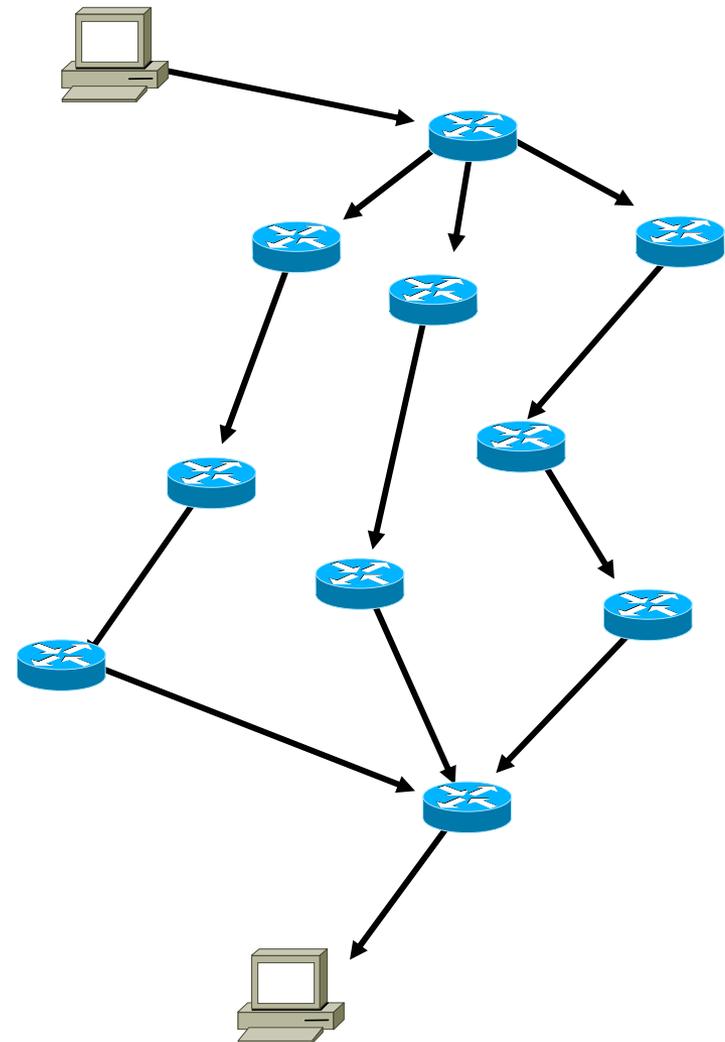
- Use Planet-Lab network as the basic test fabric
- Install two DARPA controlled test sites on Planet-Lab
- Conduct tests over Planet-Lab using performer sites, DARPA controlled sites, and up to ten other Planet-Lab sites
- Simulate tactical connections at DARPA controlled sites
  - Packet loss rates of 0% to 30%
  - Delay rates of 100 milliseconds to 3 seconds (fixed)
  - Connectivity (link) losses of 10-30 seconds every 5-30 minutes
- Weeklong test of network throughput and performance under various connectivity constraints





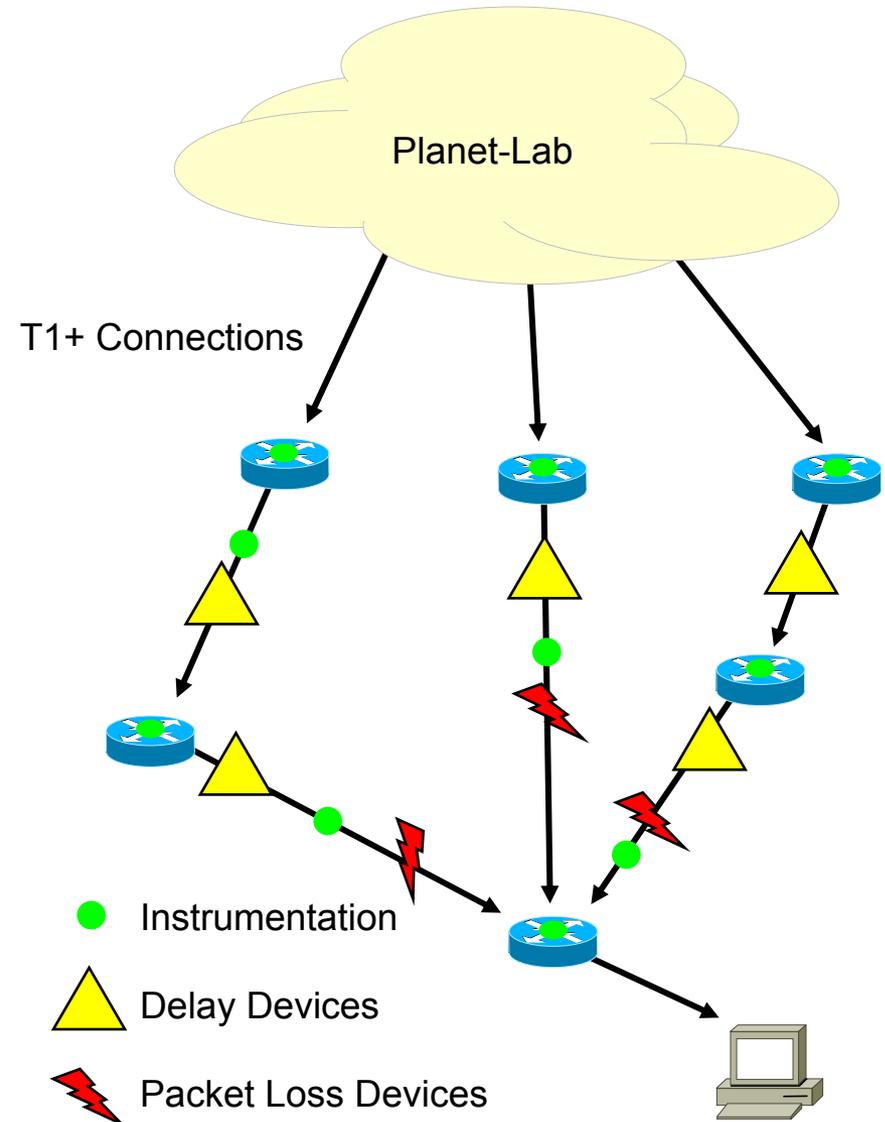
## Straight Connectivity Test

- Transfer many files of sizes varying from 10KB to 1GB, 10+ GB total
- “Pull” test—initiated from DARPA test site
- Measure with and without Control Plane improvements
- Measure with and without multiple connections



## DARPA Test Site Configuration

- Topology provided to performers
- Packet loss and delay ranges provided to performers
- Packet loss varies during test
- Delay not announced, but stays stable during test





# Performance Testing (5)

## *Milestones Two and Three*



- Milestone Two
  - Repeat Milestone One tests for improvements
  - Connect to, test, and document improvement with theater level deployable communications equipment
- Final Field Testing
  - Connect to live command posts to test and document improvement



# Cost Metrics



- Measurable against the Planet-Lab costs
  - Hardware or software additions
  - Hardware or software modifications
  - Software configuration changes and maintenance
  - Long-term training costs
- Hire an experienced, independent auditing firm
  - Work with performers during program
  - Understand work throughout
  - Assess costs at milestone testing
  - Provide independent report



# Technical Challenges



- Lessons to remember:
  - “Optimal” networking solutions that promise to always work ... don’t
  - Stick with the Internet’s “best effort” paradigm
  - Minimize infrastructure change
  - Network computing power is at the edge, use it

- DARPA Hard Problems

- Paradigm shift of how the network is used
- Making non-deterministic routing / transmissions more deterministic
- Achieving high degree of trust between devices
- Scaling to large numbers (1,000,000 x N) while doing all this
- Minimizing infrastructure changes (traffic load and hardware)

Why it will work

Must be simple enough to actually implement



# Technical Approaches



- Optimizing Network Use
- Authoritatively identify hosts to the network infrastructure
- Providing alternate traffic paths
- Filters: Hosts instruct network infrastructure



# Optimizing Network Use

*Because bandwidth does matter*



## DARPA Hard Problems

- Develop Network Infrastructure Query Protocol
  - Provide the network's current state to any authorized host
  - Cannot overload the network with requests

**Benefit: Structured mechanism to request and retrieve infrastructure information**

- Traffic Optimization Algorithms
  - Given the network state, optimize the transmissions

**Benefit: Transmit more information with less network load**

“I want the best picture you can send of downtown Baghdad ... and I need it in five minutes.”

**IMPOSSIBLE**



**Which picture to send when network endpoints have no network information?**



# Authoritative Host Identification

*How can you trust host requests?*

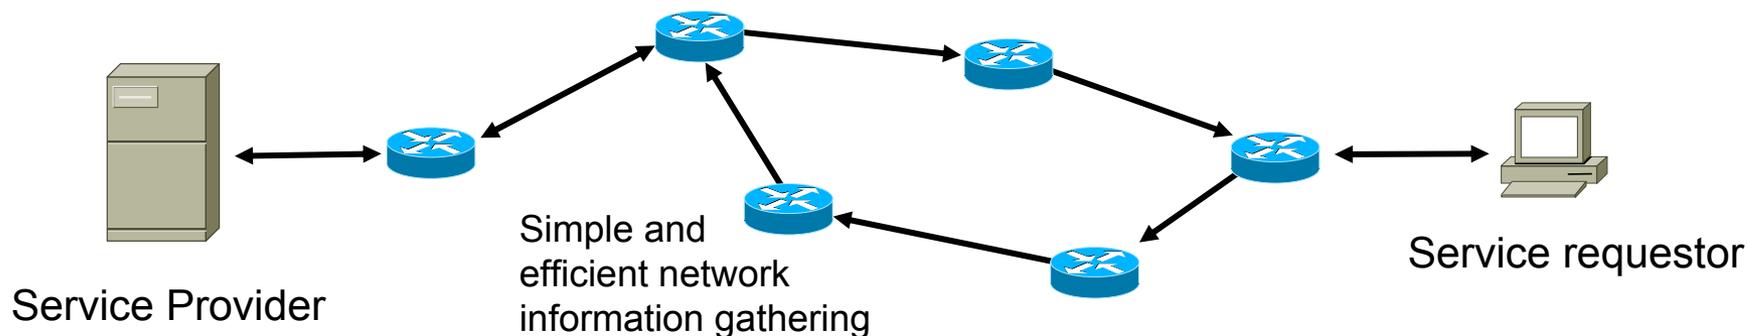


## DARPA Hard Problems

- Designing a single packet containing everything required to authenticate a request
  - Requires no additional lookup to verify authenticity
- Scalable and effective system to issue authentication tokens to make the request packets
  - Scales to 1,000,000 x N devices
  - Key management is crucial

**Benefit: Provides host level ID system**

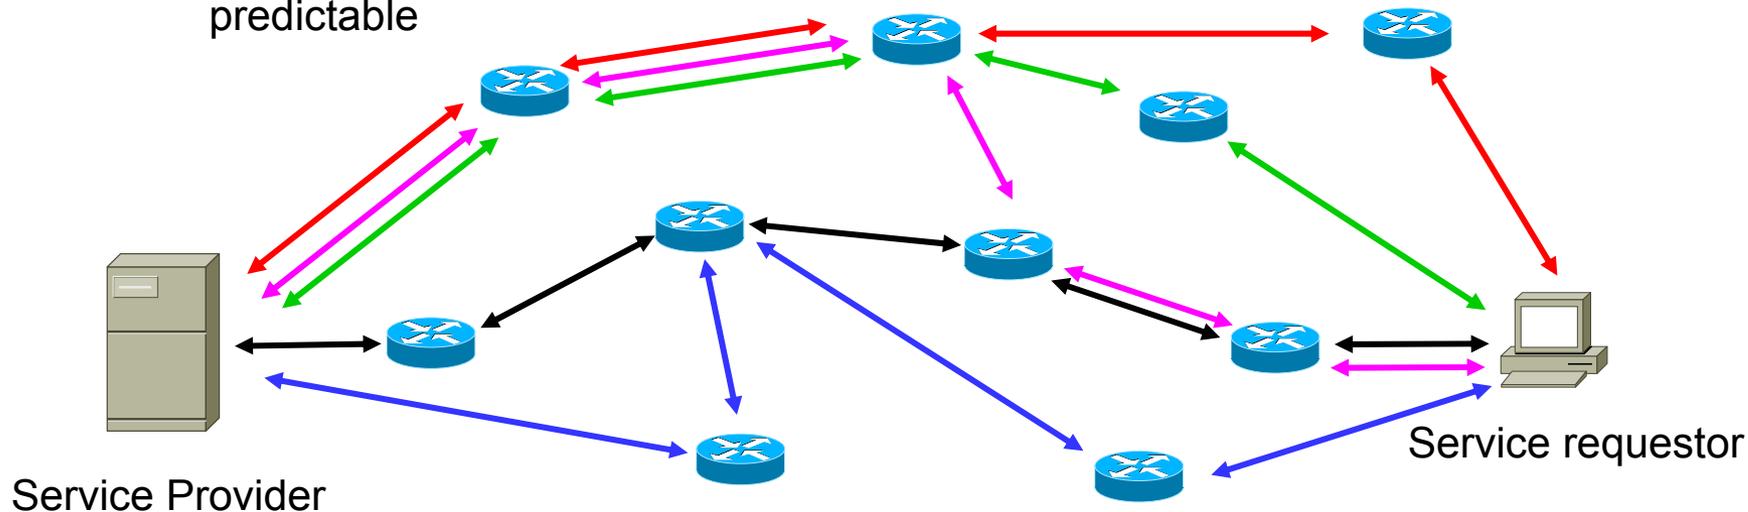
**Benefit: Usable on both DoD and commercial systems**



## DARPA Hard Problems

- Manipulate routing and route advertisement system to provide multiple paths between two points
  - Source routing not allowed
  - Route tables on routers not changeable by hosts
- Model routing system in enough detail to make the system predictable

**Benefit: Improved efficiency and throughput via multiple network paths with different characteristics between two points.**





# Infrastructure Filters and Filter Relays

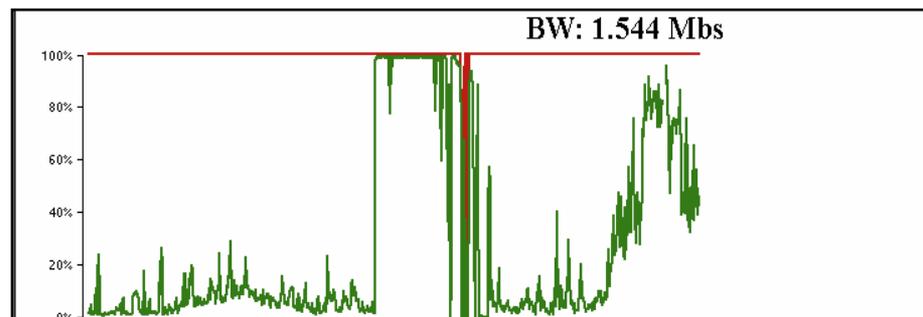
*Stretch Goal: Eliminating unwanted traffic*



## DARPA Hard Problems

- Scaling system to large numbers of hosts
- Cannot overload infrastructure capacity
- Aggregating filter requests into simpler or fewer filters
- How to pass information “upstream” through infrastructure

**Benefits: Stop attacks with existing infrastructure**



**T3 (45Mbps) Internet Access Point**



- Attack made by 200+ IPs
- Attackers may have monitored attack
  - Attack gradually slows after corrective action
- Corrective actions--upstream blocking--fixed the problem



# Military Utility and Transition



- More throughput with existing capacity
- Leverages existing network systems
  - ✓ Reinforce the Internet's "best effort" model
  - ✓ Assumes smart applications and intelligent hosts
    - ✓ Computing structure (host versus infrastructure) stays the same
  - ✓ Network infrastructure requires no significant changes
  - ✓ Control plane itself does not need to "learn" about applications
  - ✓ Can be phased in gradually (IPv4 and IPv6)
- Push unwanted traffic off chosen best path
- Transition:
  - Should be commercially available for DoD purchase
  - May be implemented in high volume service providers nodes (e.g., NIMA)

## Tasks

### Network Predictability & Efficiency

- Network Infrastructure Query Protocol
- Traffic Optimization Algorithms
- Authoritative Host Identification
- Key Management Scaling
- Application plug-ins
- Router functionality

### Alternate Traffic Paths

- DNS and BGP experimentation
- Multiple Route Experimentation
- Protocol Manipulation
- Scaling
- Applications

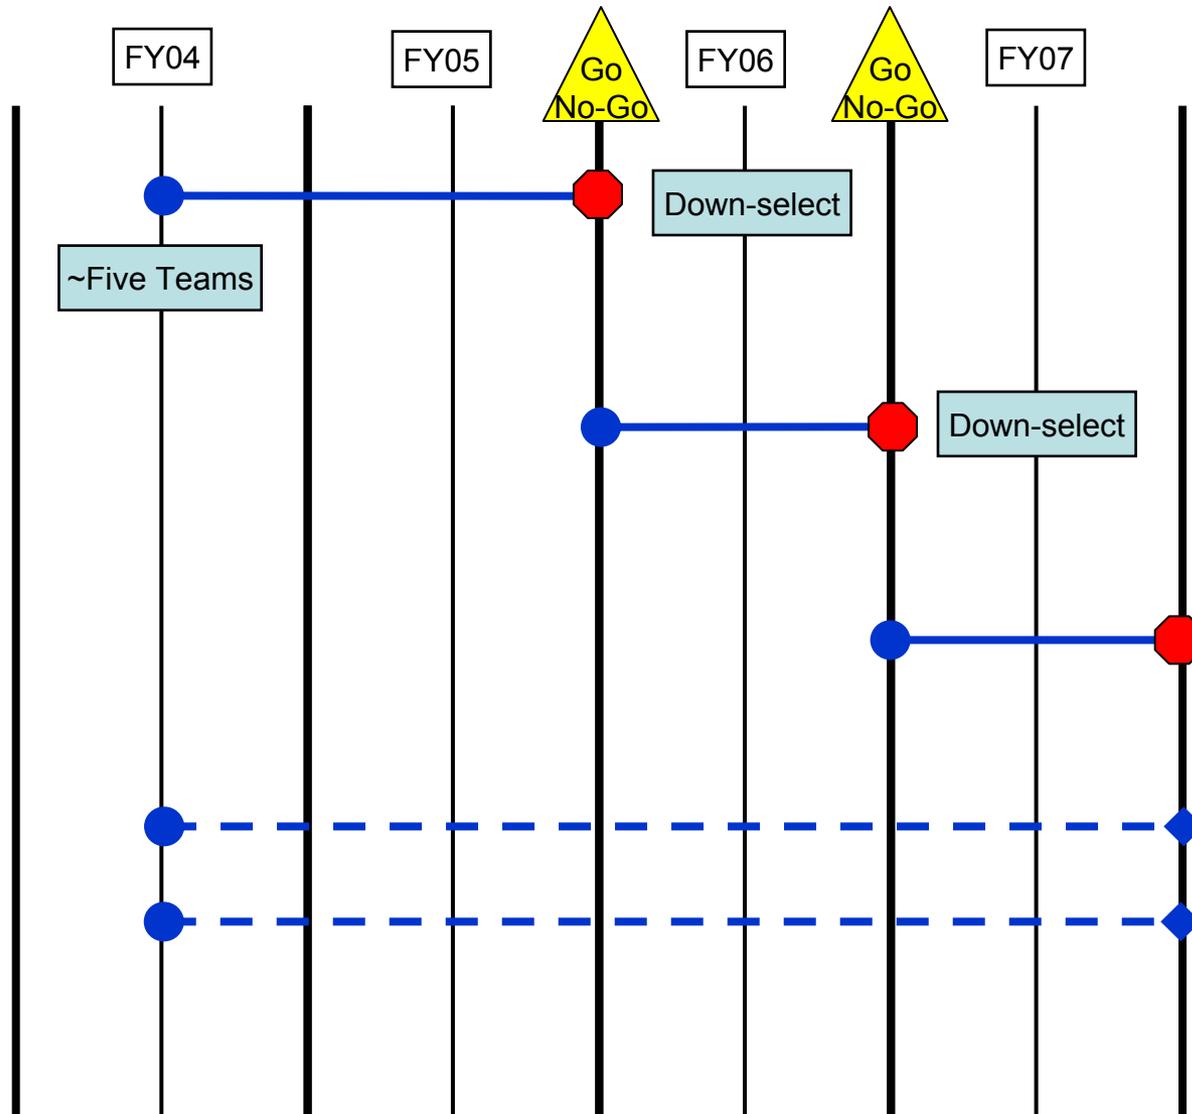
### Infrastructure Filters and Filter Relays

- Filter acceptance protocol
- Filter relay mechanism
- Scaling and optimization

### Protocol Modeling & Simulation

- Simulation Standardization
- Traffic Models and generation
- Protocol Verification

Data Collection, Testing, Verification, & Red Teaming





# Contact Information



Colonel Tim Gibson  
DARPA Advanced Technology Office  
3701 North Fairfax Drive  
Arlington, VA 22204  
Tel 703.526.4764  
Fax 703.516.7369  
tgibson@darpa.mil